



22883

PATENT TRADEMARK OFFICE

103.1032.02

This application is submitted in the name of the following inventor(s):

<u>Inventor</u>	<u>Citizenship</u>	<u>Residence City and State</u>
English, Robert M.	United States	Menlo Park, California

The assignee is Network Appliance, Inc., a corporation having an office at 495 East Java Drive, Sunnyvale, California, 94089.

Title of the Invention

Low-Overhead Threads in A High-Concurrency System

Background of the Invention

This application claims the benefit of U.S. Provisional Application No. 60/195,732, filed 4/7/00 (Attorney Docket number 103.1032.01).

*1. Field of the Invention*

This invention relates to low-overhead threads in a high-concurrency system, such as for a networked cache or file server.

2. *Related Art*

In many computing systems, it is desirable in certain circumstances to be able to process, relatively simultaneously (such as in parallel), a relatively large number of similar tasks. For example, the same or similar tasks could be performed by a server device (such as a file server) in response to requests by a number of client devices. One such circumstance is in a networked cache or file server, which maintains and processes a relatively large number of sequences of requests (sometimes called "connections"), so as to couple an information requester (such as a web client) to one or more information providers, which are also coupled to the same internetworking system. One known method in which an individual processor or a multiprocessor system is able to maintain a high degree of concurrency is for the system to process each connection using a separate processing thread. A "thread" is a locus of control within a process, indicating a spot within that process that the processor is then currently executing. In general, a thread has a relatively small amount of state information associated therewith, generally consisting only of a calling stack and a relatively small number of local variables.

High concurrency systems, such as networked caches and file servers used in an internetworking system, must generally maintain a large number of threads. Each information requester has its own separate connection for which the network cache or file server must maintain some amount of state information. Each such separate connection requires only a small amount of state information, such as approximately 100 to 200 bytes

1 of information. Since there are in many cases a relatively large number of individual  
2 connections, it would be desirable to be able to maintain state information about each  
3 such connection using only a relatively minimal amount of memory and processor over-  
4 head, while simultaneously maintaining both relatively reliable programmability and rela-  
5 tively high processing speed.

6  
7 One problem with known systems is that allocation of state information for  
8 individual threads does not generally scale well. One of the problems with relatively  
9 large numbers of individual threads is that of allocating memory space for a calling stack  
10 for each one of those threads. In a first set of known systems, stack space for individual  
11 threads is allocated statically; this has the drawback that relatively large numbers of  
12 threads require a relatively large amount of memory to maintain all such stack spaces.  
13 Although the amount of stack space statically allocated for each individual thread can be  
14 reduced significantly, this has the drawback that operations that can be performed by each  
15 individual thread are similarly significantly restricted. In a second set of known systems,  
16 stack space for individual threads is allocated dynamically; this has the drawback that the  
17 minimum size for dynamic allocation of memory is generally measured in kilobytes, re-  
18 sulting in substantial unnecessary memory overhead. Although virtual memory can be  
19 used to store and retrieve stack space for individual threads in smaller increments, this has  
20 the drawback that compression and decompression of stack space for individual threads  
21 imposes substantial unnecessary processor overhead. In a third set of known systems,  
22 such as those using the Java programming language, dynamic memory allocation is used

1 to store and retrieve stack space for individual threads; this has the drawback that each  
2 procedure call within each thread imposes substantial unnecessary processor overhead.

3  
4 An additional problem is introduced by the particular use made of multi-  
5 threading by the WAFL file system (as described in the Incorporated Disclosures). In the  
6 WAFL file system, the C language "setjmp" and "longjmp" routines are combined with  
7 message passing among threads so as to support high concurrency using threads. In par-  
8 ticular, the requester of an initial file request to the WAFL file system packages the re-  
9 quest in a message, which the WAFL file system processes using ordinary procedural  
10 program code, so long as data is available for processing the request and the thread need  
11 not have its execution suspended. If the thread is suspended for any reason (such as if a  
12 resource is not available,) the WAFL file system: (1) requests the needed resource, (2)  
13 queues the message for signaling when the resource is available, and (3) calls the C rout-  
14 ing "longjmp" to return to the origin of the routine for processing the message. Thus, the  
15 WAFL file system restarts processing the entire message from the very beginning until all  
16 needed resources are available and processing can complete without suspension. While  
17 this use of multithreading by the WAFL file system has the advantage that programmers  
18 do not need to encode program state when a routine is suspended, it has the disadvantage,  
19 when combined with multithreading, that all necessary data structures (to process any ar-  
20 bitrary message) must be collected before the entire message can be processed. In an in-  
21 ternetworking environment, collecting all such structures can be difficult and subject to  
22 error.

1  
2 Accordingly, it would be advantageous to provide a technique for creating  
3 and using relatively low-overhead threads in a high-concurrency system, such as for a  
4 networked cache or file server, that is not subject to drawbacks of the known art.  
5

#### 6 Summary of the Invention

7

8 The invention provides a method and system for providing the functionality  
9 of dynamically-allocated threads in a multithreaded system in which the operating system  
10 provides only statically-allocated threads. With this functionality, a relatively large num-  
11 ber of threads can be maintained without a relatively large amount of overhead (either in  
12 memory or processor time), and it remains possible to produce program code without un-  
13 due complexity.  
14

15 In a preferred embodiment, a plurality of dynamically-allocated threads are  
16 simulated using a single statically-allocated thread, but with state information regarding  
17 each dynamically-allocated thread maintained within the single statically-allocated thread.  
18 The single statically-allocated thread includes, for each procedure call that would other-  
19 wise introduce a new dynamically-allocated thread, a memory block including: (1) a rela-  
20 tively small procedure call stack for the new dynamically-allocated thread, and (2) a rela-  
21 tively small collection of local variables and other state information for the new dynami-  
22 cally-allocated thread. When using multithreading in the WAFL file system, high

1 concurrency among threads can be maintained without any particular requirement that the  
2 program code maintain a substantial amount of state information regarding each dynami-  
3 cally-allocated thread. Each routine in the WAFL file system that expects to be sus-  
4 pended or interrupted need maintain only a collection of entry points into which the rou-  
5 tine is re-entered when the suspension or interruption is completed. A feature of the C  
6 language preprocessor allows the programmer to generate each of these entry points  
7 without substantial additional programming work, with the aid of one or more program-  
8 ming macros.

9  
10 The invention provides an enabling technology for a wide variety of appli-  
11 cations for multithreaded systems so as to obtain substantial advantages and capabilities  
12 that are novel and non-obvious in view of the known art. Examples described below pri-  
13 marily relate to networked caches and file servers, but the invention is broadly applicable  
14 to many different types of automated software systems.

#### 15 16 Brief Description of the Drawings

17  
18 Figure 1 shows a block diagram of a system for providing functionality of  
19 low-overhead threads in a high-concurrency system, such as for a networked cache or file  
20 server.

Figure 2 shows a process flow diagram of a system for providing functionality of low-overhead threads in a high-concurrency system, such as for a networked cache or file server.

### Detailed Description of the Preferred Embodiment

In the following description, a preferred embodiment of the invention is described with regard to preferred process steps and data structures. Embodiments of the invention can be implemented using general-purpose processors or special purpose processors operating under program control, or other circuits, adapted to particular process steps and data structures described herein. Implementation of the process steps and data structures described herein would not require undue experimentation or further invention.

### *Lexicography*

The following terms refer or relate to aspects of the invention as described below. The descriptions of general meanings of these terms are not intended to be limiting, only illustrative.

- **client and server** — In general, these terms refer to a relationship between two devices, particularly to their relationship as client and server, not necessarily to any particular physical devices.

1  
2 For example, but without limitation, a particular client device in a first relationship  
3 with a first server device, can serve as a server device in a second relationship with  
4 a second client device. In a preferred embodiment, there are generally a relatively  
5 small number of server devices servicing a relatively larger number of client de-  
6 vices.

7  
8 ○ **client device and server device** — In general, these terms refer to devices taking  
9 on the role of a client device or a server device in a client-server relationship (such  
10 as an HTTP web client and web server). There is no particular requirement that  
11 any client devices or server devices must be individual physical devices. They can  
12 each be a single device, a set of cooperating devices, a portion of a device, or some  
13 combination thereof.

14  
15 For example, but without limitation, the client device and the server device in a  
16 client-server relation can actually be the same physical device, with a first set of  
17 software elements serving to perform client functions and a second set of software  
18 elements serving to perform server functions

19  
20 As noted above, these descriptions of general meanings of these terms are  
21 not intended to be limiting, only illustrative. Other and further applications of the inven-  
22 tion, including extensions of these terms and concepts, would be clear to those of ordinary



1 skill in the art after perusing this application. These other and further applications are  
2 part of the scope and spirit of the invention, and would be clear to those of ordinary skill  
3 in the art, without further invention or undue experimentation.

#### 4 5 *System Elements*

6  
7 Figure 1 shows a block diagram of a system for providing functionality of  
8 low-overhead threads in a high-concurrency system, such as for a networked cache or file  
9 server.

10  
11 A system 100 includes a networked cache or file server (or other device)  
12 110, a sequence of input request messages 120, and a set of software elements 130.

13  
14 The networked cache or file server (or other device) 110 includes a com-  
15 puter having a processor, program and data memory, mass storage, a presentation ele-  
16 ment, and an input element, and is coupled to a communication network. As used herein,  
17 the term "computer" is intended in its broadest sense, and includes any device having a  
18 programmable processor or otherwise falling within the generalized Turing machine  
19 paradigm. The mass storage can include any device for storing relatively large amounts  
20 of information, such as magnetic disks or tapes, optical devices, magneto-optical devices,  
21 or other types of mass storage.

1           The input request messages 120 include a set of messages requesting the  
2 networked cache or file server 110 to perform actions in response thereto. In a preferred  
3 embodiment, the actions to be performed by the networked cache or file server 110 will  
4 involve access to the mass storage or to the communication network. In a preferred em-  
5 bodiment, the input request messages 120 are formatted in a known request protocol, such  
6 as NFS, CIFS, HTTP (or variants thereof), but there is no particular requirement for the  
7 input request messages 120 to use these known request protocols or any other known re-  
8 quest protocols. In a preferred embodiment, the networked cache or file server 110 re-  
9 sponds to the input request messages 120 with both: (1) a condign set of responsive ac-  
10 tions involving the mass storage or the vacation network, and (2) a condign response to  
11 the input request messages 120, the response to the input request messages 120 preferably  
12 taking the form of a set of response messages (not shown.)

13  
14           The software elements 130 include a set of programmed routines to be per-  
15 formed by the networked cache or file server 110, using the functionality of low-overhead  
16 threads and high-concurrency as described herein. Although particular program code is  
17 described herein with regard to the programmed routines, there is no particular reason that  
18 the software elements 130 must use the specific program code described herein, or any  
19 other specific program code.

## Method of Operation

Figure 2 shows a process flow diagram of a system for providing functionality of low-overhead threads in a high-concurrency system, such as for a networked cache or file server.

A method 200 includes a set of flow points and a set of steps. The system 100 performs the method 200. Although the method 200 is described serially, the steps of the method 200 can be performed by separate elements in conjunction or in parallel, whether asynchronously, in a pipelined manner, or otherwise. There is no particular requirement that the method 200 be performed in the same order in which this description lists the steps, except where so indicated.

At a flow point 210, the networked cache or file server 110 is ready to receive and respond to the input request messages 120.

At a step 211, the networked cache or file server 110 receives an input request message 120, and forwards that input request message 120 to an appropriate software element 130 for processing. In a preferred embodiment, the step 211 includes performing a calling sequence for the software element 130, including possibly creating a simulated dynamically allocated thread (that is, a thread simulated so as to appear to be dynamically-allocated, hereinafter sometimes called a “simulated thread” or an “S-

thread") within which the software element 130 is performed. Thus, the software element 130 can be created using program code that assumes that the software element 130 is performed by a separate thread and does not demand relatively excessive resources (either memory or processor time.)

As part of step 211, the networked cache or file server 110 allocates a procedure call block 131 and a local variable block 132, for use by the simulated dynamically-allocated thread performed by the software element 130. The procedure call block 131 includes a set of input variables for input to the software element 130, a set of output variables for output from the software element 130, and such other stack element as is known in the art of calling stacks for procedure calls. The local variable block 132 includes a set of locations in which to store local variables for the software element 130.

As part of step 211, the networked cache or file server 110 determines whether the software element 130 is a subroutine of a previously called software element 130 in the same simulated thread. If so, the networked cache or file server 110 indicates that fact in a block header 133 for the software element 130, so as to point back to the particular software element 130 that was the parent (calling) software element 130. If not, the networked cache or file server 110 does not indicate that fact in the block call or block header for the software element 130.

1           As part of this step, the networked cache or file server 110 determines  
2 whether the software element 130 is to be performed by a new simulated thread. If so, the  
3 networked cache or file server 110 adds the new thread block 134 to a linked list 135 of  
4 thread blocks 134 to be performed in turn according to a scheduler. In a preferred em-  
5 bodiment, the scheduler simply performs each simulated thread corresponding to the next  
6 thread block 134 in round-robin sequence, so that each simulated thread corresponding to  
7 a thread block 134 is performed in its turn, until it is suspended or completes. However,  
8 in alternative embodiments, the scheduler may select simulated threads in other than a  
9 round-robin sequence, so as to achieve a desired measure of quality of service, or other  
10 administrative goals.

11           At a step 212, the networked cache or file server 110 chooses the simulated  
12 thread for execution. The simulated thread, with appropriate data completed for the pro-  
13 cedure call block 131 and local variable block 132, is performed in its turn, until it is sus-  
14 pended or completes. If the simulated thread is capable of completing its operation with-  
15 out being suspended or interrupted, the scheduler selects the next thread block 134 in the  
16 linked list of thread blocks 134 to be performed in turn.

17  
18  
19           After this step, the method 200 has performed one round of receiving and  
20 responding to input request messages 120, and is ready to perform another such round so  
21 as to continuously receive and respond to input request messages 120.

1 The method 200 is performed one or more times starting from the flow  
2 point 210 and continuing therefrom. In a preferred embodiment, the networked cache or  
3 file server 110 repeatedly performs the method 200, starting from the flow point 210 and  
4 continuing therefrom, so as to receive and respond to input request messages 120 periodically  
5 cally and continuously.

### 6 7 *Program Structures*

8  
9 A set of program structures in a system for providing functionality of low-  
10 overhead threads in a high-concurrency system, such as for a networked cache or file  
11 server, includes one or more of, or some combination of, the following:

- 12  
13 ○ A set of program structures for declaring and creating a dynamically-allocated thread  
14 in a system in which threads are usually statically-allocated;

```
15  
16 typedef struct {  
17     // local variables  
18     int arg;    // an example, not necessary  
19     } function_msg;  
20
```

21 In the program structure above, the definition for the structure type “func-  
22 tion\_msg” includes: (1) the local variables for the dynamically-allocated thread, (2) any

input arguments to the dynamically-allocated thread, in this case just the one variable “arg”, and (3) any output arguments from the dynamically-allocated thread, in this case none.

◦ A set of program structures for denoting program code entry-points for a simulated thread;

```
static void
function_sthread(sthread_msg *m)
{
    function_msg * const msg = m->data;

    STHREAD_START_BLOCK (m);
    // executable C code
    STHREAD_RESTART_POINT (m);          // an example
blocking point
    // executable C code
    STHREAD_COND_WAIT (m, cond (m)); // encapsulated
blocking point
    // executable C code
    STHREAD_END_BLOCK;
    free (msg);
}
```

The program structure above includes, in its definition for the function “function\_sthread”, an initial program statement obtaining access to the local variables for the simulated thread. This is the statement referring to “m -> data”.

1  
2           The program structure above includes a definition for a start-point for the  
3 simulated thread. This is the statement "STHREAD\_START\_BLOCK (m)", which  
4 makes use of a macro defined for the name "STHREAD\_START\_BLOCK".

5  
6           The program structure above includes a definition for a restart-point for the  
7 simulated thread. This is the statement "STHREAD\_RESTART\_POINT (m)", which  
8 makes use of a macro defined for the name "STHREAD\_RESTART\_POINT".

9  
10           The program structure above includes a definition for a conditional-wait  
11 point (a possible suspension of the simulated thread) for the simulated thread. This is the  
12 statement "STHREAD\_COND\_WAIT(m, cond(m))", which makes use of a macro de-  
13 fined for the name "STHREAD\_COND\_WAIT".

14  
15           The program structure above includes, in its definition for the function  
16 "function\_sthread", a closing program statement for ending the simulated thread. This  
17 is the statement "STHREAD\_END\_BLOCK", which makes use of a macro defined for  
18 the name "STHREAD\_END\_BLOCK". The program structure above also includes a  
19 statement for freeing any data structures used by the simulated thread. This is the state-  
20 ment "free(msg)".



The macro definitions for “STHREAD\_START\_BLOCK”, “STHREAD\_RESTART\_POINT”, and “STHREAD\_END\_BLOCK” collectively form a C language “case” statement.

- The macro “STHREAD\_START\_BLOCK” includes the preamble to the “case” statement:

```
#define STHREAD_START_BLOCK (m) switch (m -> line) { case 0:
```

- The macro “STHREAD\_RESTART\_POINT” includes an intermediate restart point in the “case” statement:

```
#define STHREAD_RESTART_POINT(m) case __LINE__: m -> line  
= __LINE__
```

The restart point uses the C preprocessor to generate tags that the switch statement uses as branch points. The C macro `__LINE__` substitutes the line number of the file being processed, so a series of restart points generates a series of unique cases within the switch. Setting `m -> line` to the case just entered means that if the procedure is re-entered the switch statement will branch to the restart point and continue.

- The macro “STHREAD\_START\_BLOCK” includes the close of the “case” statement:

```
#define STHREAD_END_BLOCK }
```

Thus, the C preprocessor generates a “case” statement in response to use of these macros, which allows the programmer to easily specify each of the proper restart points of the routine.

- A set of program structures for suspending and restarting simulated threads;

```
#define STHREAD_COND_WAIT(m, c) \
STHREAD_RESTART_POINT(m); \
{ \ if (c) \
    sthread_suspend(); \
}
```

At an individual restart point, the programmer can use the macro “STHREAD\_COND\_WAIT” to conditionally either wait for an operation to complete, or to suspend and restart the simulated thread while waiting for resources for the operation to complete.

- A set of program structures for initiating simulated threads;

- The macro “STHREAD\_INIT” allocates memory for the simulated thread, sets the C preprocessor value `__LINE__` to zero, sets the value of “data” to the private stack area of the particular simulated thread, and sets a value for “handler” to a function passed to the macro as an argument.

```
#define STHREAD_INIT(m, msg, handler) \ m = malloc(sizeof(*m)); \
msg = zalloc(sizeof(*msg)); \ m -> line = 0; \ m -> data = msg; \ m ->
handler = handler
```

- A set of program structures for actually performing the simulated thread;

```
void
function(int arg)
{
    function_msg *msg;
    sthread_msg *m;

    STHREAD_INIT(m, msg, function_sthread);
    msg->arg = arg;

    sthread_run(m);
}
```

The program structure above includes, in its definition for the function “function”, program code for creating the data blocks for the simulated thread, and for placing data in those data blocks. These are the statements “STHREAD\_INIT(m, msg, function\_thread)” and “msg -> arg = arg”, which make use of a macro defined for the name “STHREAD\_INIT”.

- A set of program structures for scheduling performance of simulated threads;

```

switch (m->line) {    // a field in sthread_msg
case 0:
    // executable C code
    STHREAD_RESTART_POINT(m);
    // executable C code
    STHREAD_RESTART_POINT(m);
    // executable C code
}

```

The program structure above includes, in its definition for the function “function”, program code for creating the data blocks for the simulated thread, and for placing data in those data blocks. These are the statements “STHREAD\_INIT(m, msg, function\_thread)” and “msg -> arg = arg”, which make use of a macro defined for the name “STHREAD\_INIT”.

- 1   ○ A set of program structures for suspending and resuming performance of simulated
- 2   threads.

```

4      typedef struct sthread_msg {
5          int line;
6          void *data;
7          void (*handler)(sthread_msg *);
8      }
9      jmp_buf sthread_env;
10     void
11     sthread_run(sthread_msg *m)
12     {
13         if (!setjmp(sthread_env)) {
14             m->handler(m);
15             free(m);
16         }
17     }
18     void
19     sthread_suspend()
20     {
21         longjmp(sthread_env, 0);
22     }
23     sthread_msg *suspended_sthread;
24     int ready;
25     int
26     cond(sthread_msg *m)
27     {
28         if (ready)

```

```
1      return 1;
2      suspended_sthread = m;
3      sthread_suspend();
4  }
5
6
7  int
8  set_cond()
9  {
10     ready = 1;
11     if (suspended_sthread) {
12         sthread_msg *m = suspended_sthread;
13         suspended_sthread = 0;
14         sthread_run(m);
15     }
16 }
17 // cond() changed
18 sthread_run(suspended_sthread);
```

and

- A set of program structures for performing simulated threads in conjunction with the WAFL file system, as shown above.

### *Generality of the Invention*

The invention has general applicability to various fields of use, not necessarily related to the services described above. For example, these fields of use can include devices other than file servers.

Other and further applications of the invention in its most general form, will be clear to those skilled in the art after perusal of this application, and are within the scope and spirit of the invention.

### *Technical Appendix*

The technical appendix enclosed with this application is hereby incorporated by reference as if fully set forth herein, and forms a part of the disclosure of the invention and its preferred embodiments.

### *Alternative Embodiments*

Although preferred embodiments are disclosed herein, many variations are possible which remain within the concept, scope, and spirit of the invention, and these variations would become clear to those skilled in the art after perusal of this application.